

Breakout: Terabit Network End Systems Agenda and Questions

Day 1: Breakout Session 1:

- Welcome and session logistics - Don Petravick and Bill Allcock
- **Opening presentation** – Advances in storage systems technologies and I/O subsystems – Gary Grider
- Discussions on questions 1, 2, and 3:
- Summary of discussions - scribe

Day 1: Breakout Session 2:

- Welcome and session logistics – Don Petravick and Bill Allcock
- **Case study presentation:** Network Optimized storage/file system in leadership computing environment – Galen Shipman
- Continued discussion on questions 1.2, and 3
- Summary of discussions - scribe

Day 2: Breakout Session 1

- Welcome and session logistics – Don Petravick and Bill Allcock
- **Case study presentation** – Campus/Site networks environment -
- Phil DeMar
- Discussions on questions 4,5, and 6:
- Summary of discussions - scribe

Day 2: Breakout session 2:

- Welcome and session logistics – Don Petravick and Bill Allcock
 - **Case study presentation:** Science data center: Cloud computing, cluster computing , and high-end data servers – Brent Draney
 - Applications test case.
 - Continued discussions on Question 4,5, and 6
 - Summary of discussions - scribe
-

Session 1: Terabit Network End Systems

Context:

- High-performance storage systems for sourcing and sinking massive data sets over the emerging 100 Gbps links
- Network-aware optimized file systems and metadata
- Host/system provisioning
- Parallel I/O and file systems
- Data transfer at scale

1. **High Level Use Case Development and Testbeds** : How will Terabit network systems be used? What are the key characteristics of each? How could we construct a community testbed that could serve all these use cases? As we discuss the remaining questions we should consider the answers for each of the use cases.
2. **Host/end systems** – Host systems and the associated network stacks and interfaces performance has emerge as a major obstacle to high-speed data transfer, especially at 10 Gbps. The obstacle will likely worsen at 100 Gbps and Tbps speeds. What will be the major challenges that must be resolved in order to address these host/systems issues? Some talking points include:
 1. Local Area Network environments
 2. Host network stack and transport protocols
 3. Data transfer protocols.
 4. Disk/storage performance
3. **File systems and metadata** – File systems and metadata are important components in that if they are properly designed they can enhance I/O performance to scaling beyond the limits of traditional storage technology. What are challenges in developing I/O subsystems for 100x performance scaling and beyond? What are the associated challenges in harnessing parallel file systems features and I/O concurrencies in to improve data transfer performance in host/storage systems at terascale?
 1. Parallel file systems
 2. I/O concurrency
 3. Metadata scaling
 4. SAN protocol over WAN transports
 5. Data streaming via parallel TCP

4. **Data centers, data grid, data cluster, host provisioning and performance monitoring for high-speed data transfer** – What are the technical issues in adapting commercial data center approaches to science data storage, management, and distribution? How should end host systems be provisioned to take full advantage of the abundant bandwidth available in the backbone? What are the challenges of monitoring end systems to achieve complete end-to-end monitoring and fault diagnosis? Talking points:
 1. Differentiated services for end systems
 2. Integrated backbone performance monitoring with end host monitoring
 3. Interfacing end host systems to backbone networks

5. **TCP/UDP transport protocols and viable transport protocols for massive data transfers** – Is TCP and its variants viable transport protocols for data transmission at Tbps? Given the Given that it requires heroic and coordinated effort for TCP to deliver line rate performance, what should be done to improve the performance at 100 Gbps and beyond. Given its wide deployment, is a TCP alternative of TCP worth considering? Some talking points
 1. RDMA over Provisioned WAN circuits, SAN (FB, FC, SCSSI, etc.) protocols over provisioned WAN circuits
 2. Composable/adaptable transport protocol
 3. Others

6. **Storage/File Systems/Hosts Testbeds** – The next-generation of terabits end systems may be radically different from existing ones. Is there a compelling need for experimental/testbed environment that will facilitate the prototyping and testing end systems hardware, software, and protocols? What will be relation of this testbed to wide area terabits network testbeds creating a comprehensive experimental environment for developing end-to-end terabit network technologies? How can DOE leverage its exascale simulation efforts to augment testbed activities? Discussion issues
 1. Joint DOE/industry partnerships on next-generation end systems
 2. End-to-end testbeds
 3. Leveraging exascale hardware approaches such as co-design/simulations
 4. I/O bandwidth modeling and analysis for exascale

7. **Community – Host systems provisioning, virtualization, and performance monitoring** – DOE's scientists typically collaborate and work large groups (Climate -

ESG, HEP-OSG, etc.) to undertake complex science effort. What are the challenges in developing and deploying and managing automated community-based advanced services? What is the role network virtualization in this effort? Issues for considerations:

1. community-based danced network services
2. Users' level network virtualization services
3. Community-based access control credential management
4. others

8. Other Questions